


Future Technology Challenges for Collaboration: Customizing the Grid

Jon B. Weissman


**Distributed Computing Systems Group
Department of Computer Science
University of Minnesota**

*DOE National Collaboratories Meeting
ANL, August 10-12, 2004*



Outline

- Enabling=>Enticing collaboration
- It's the software, stupid
- Our vision: customized Grid services
- Research challenges




Enticing Collaboration

- Many elements are needed for a successful collaboration
 - incentive
 - well-defined boundaries and separation
 - perception of value-added
 - **substrate** to promote interaction
- The Grid is a place where collaboration can happen
 - temporally and spatially
 - software is an important conduit for Grid collaboration




- once groups decide to collaborate ... make it easier ...
- make it easy to deploy services dynamically
- allow services to adapt to Grid dynamics






Tomorrow: Enticing Collaboration

- Software-based collaboration
 - group A provides software A' (computation)
 - group B provides software B' (visualization)
- Problem 1
 - the coupling of A' and B' require standard-interfaces to enable A' and B' to evolve independently
 - solution: Grid services!
- Problem 2
 - B' is not exactly what A needs and A' is not exactly what B needs



Obstacles


- Group A decides it would be easier to implement B' in-house and B also decides this for A'
 - typical in “roll-your-own” cultures
 - no need to collaborate or ...
 - groups may also choose to live w/o the external services
- Group B might agree to tailor B' for A ... but this effort cannot be re-used ...
 - another group C might want B's software and again the same problem occurs
 - ad-hoc modifications of software are not durable



Blue Sky


Suppose value-added services were easy to find **and** tunable to a group's needs?

The Grid was a “place” that groups could find collaborators




Shape of a Solution


- Disciplined flexibility
- Flexible Grid service
 - standards-based
 - exposes and supports flexibility in a **uniform** way
- Customized Grid service
 - specific instance of a flexible Grid service
 - achieves client-centric notion of multi-dimensional QoS




Flexible Grid Service Vision




solver is optimized for speed only for each request




solver is optimized for throughput




solver is optimized for reliability



“Dialing for Service”




“I want performance P , reliability R , using algorithm A , with cost C ”




Control Interface

- Control interface consists of software “dials”
 - dials are “turned” to make an instance customized
 - a dial maps to a QoS feature
- Categorizing QoS features
 - performance
 - throughput X, max completion time Y, response Z, sustained rate S
 - reliability
 - X 9's of reliability, checkpointing Y/N
 - security
 - X encryption bits
 - functionality
 - algorithm A, B, C (e.g. SW, FASTA, BLAST)
 - accuracy
 - X digits
 - ...




Challenges

- How can service flexibility be implemented and what is the incentive for the service provider to provide it?
 - could the control interface be retrofitted ... how?
- How does the user or client notion of QoS map to the various dial settings?




Challenges (cont'd)

- How can competing customizations to the same service be handled
 - e.g. group A and group B want conflicting features from the same service?
 - can we guarantee some kind of “QoS” isolation?
- How can feature dependence be represented and handled?
 - e.g. turning the dial on a reliability dial may come at the cost of performance?



Challenges (cont'd)

- How can service discovery be enhanced to include customized services
 - e.g. a group may wish to locate a genomic service which provides a performance dial that can deliver a throughput between X and Y?
- Can composite services (a service that contains or uses other services) be made flexible?
 - e.g. solver uses compute and storage services
 - how to achieve end-to-end QoS?



Challenges (cont'd)

- To what extent does the dynamics and unpredictable nature of the Grid impact customization?
 - are QoS guarantees merely stochastic?
- What common middleware could be provided in support of flexible services?